

L Number	Hits	Search Text	DB	Time stamp
1	200	((power adj control) or (power\$3 adj down) or (conserv\$3 adj power)) with cache	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:01
2	1482	(disabl\$3 or turn-off or turn adj off or shut adj down) with cache	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:03
3	4031	microinstruction	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:05
4	4	((disabl\$3 or turn-off or turn adj off or shut adj down) with cache) same microinstruction	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:03
5	4708	microinstruction or micro-instruction	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:06
6	9989	(microinstruction or micro-instruction) or uop	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:06
7	5	((disabl\$3 or turn-off or turn adj off or shut adj down) with cache) same ((microinstruction or micro-instruction) or uop)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:06
8	1	(((disabl\$3 or turn-off or turn adj off or shut adj down) with cache) same ((microinstruction or micro-instruction) or uop)) not (((disabl\$3 or turn-off or turn adj off or shut adj down) with cache) same microinstruction)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:07
9	2	power with cache with ((microinstruction or micro-instruction) or uop)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:08
10	1101	711/154.cc1s.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:09
11	370	711/128.cc1s.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:10
12	318	711/136.cc1s.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:10
13	556	711/133.cc1s.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:10
14	462	711/156.cc1s.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:10
15	539	713/320.cc1s.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:10
16	434	713/322.cc1s.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:10
17	419	713/324.cc1s.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/07/11 16:11

18	9	((power adj control) or (power\$3 adj down) or (conserv\$3 adj power)) with cache) and microinstruction	USPAT; US-PGPUB; EPO; JPO; IBM TDB	2003/07/11 16:11
19	4	5781783.URPN.	USPAT	2003/07/11 16:14
20	1		USPAT	2003/07/11 16:17
21	1		USPAT	2003/07/11 16:17
22	1		USPAT	2003/07/11 16:17
23	1		USPAT	2003/07/11 16:18
24	1		USPAT	2003/07/11 16:19
25	1		USPAT	2003/07/11 16:19
26	1		USPAT	2003/07/11 16:20
27	1		USPAT	2003/07/11 16:20
28	1		USPAT	2003/07/11 16:21
29	1		USPAT	2003/07/11 16:21
30	1		USPAT	2003/07/11 16:21

Welcome to IEEE Xplore® Your search matched [0] of [951805] documents.

- [Home](#)
- [What Can I Access?](#)
- [Log-out](#)

Tables of Contents

- [Journals & Magazines](#)
- [Conference Proceedings](#)
- [Standards](#)

Search

- [By Author](#)
- [Basic](#)
- [Advanced](#)

Member Services

- [Join IEEE](#)
- [Establish IEEE Web Account](#)
- [Access the IEEE Member Digital Library](#)

[!\[\]\(b64b40baaee5acddc1eab8538ba84754_img.jpg\) Print Format](#)

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#)
[Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#)
[No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Welcome to IEEE Xplore® Your search matched **[0]** of **[950011]** documents.

- [○- Home](#)
- [○- What Can I Access?](#)
- [○- Log-out](#)

Tables of Contents

- [○- Journals & Magazines](#)
- [○- Conference Proceedings](#)
- [○- Standards](#)

Search

- [○- By Author](#)
- [○- Basic](#)
- [○- Advanced](#)

Member Services

- [○- Join IEEE](#)
- [○- Establish IEEE Web Account](#)
- [○- Access the IEEE Member Digital Library](#)

 [Print Format](#)

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#)
[Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#)
[No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Welcome to IEEE Xplore® Your search matched [0] of [950011] documents.

- [Home](#)
- [What Can I Access?](#)
- [Log-out](#)

Tables of Contents

- [Journals & Magazines](#)
- [Conference Proceedings](#)
- [Standards](#)

Search

- [By Author](#)
- [Basic](#)
- [Advanced](#)

Member Services

- [Join IEEE](#)
- [Establish IEEE Web Account](#)
- [Access the IEEE Member Digital Library](#)

[!\[\]\(7bc43b319a082987e20f7bf78f4bab80_img.jpg\) Print Format](#)

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#)
[Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#)
[No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Welcome to IEEE Xplore®

- Home
- What Can I Access?
- Log-out

Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

Search

- By Author
- Basic
- Advanced

Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

 Print FormatYour search matched **8** of **950011** documents.

A maximum of **8** results are displayed, **25** to a page, sorted by **Relevance in descending order**. You may refine your search by editing the current search expression or entering a new one the text box. Then click **Search Again**.

Results:

Journal or Magazine = **JNL** Conference = **CNF** Standard = **STD****1 RIPAC: a VLSI processor for speech recognition***Licciardi, L.; Paolini, M.; Tasso, R.; Torielli, A.; Cecinati, R.;*

Custom Integrated Circuits Conference, 1989., Proceedings of the IEEE 1989 , 15-18

May 1989

Page(s): 20.6/1 -20.6/4

[\[Abstract\]](#) [\[PDF Full-Text \(288 KB\)\]](#) **IEEE CNF****2 On the scheduling algorithm of the dynamically trace scheduled VLIW architecture***Ferreira de Souza, A.; Rounce, P.;*

Parallel and Distributed Processing Symposium, 2000. IPDPS 2000. Proceedings. 14th International , 1-5 May 2000

Page(s): 565 -572

 Print Format[\[Abstract\]](#) [\[PDF Full-Text \(172 KB\)\]](#) **IEEE CNF****3 An image signal multiprocessor on a single chip***Maruyama, M.; Nakahira, H.; Araki, T.; Sakiyama, S.; Kitao, Y.; Aono, K.; Yamada, H.;*

Solid-State Circuits, IEEE Journal of , Volume: 25 Issue: 6 , Dec. 1990

Page(s): 1476 -1483

[\[Abstract\]](#) [\[PDF Full-Text \(720 KB\)\]](#) **IEEE JNL****4 Code optimization technique for general-purpose-DSP compiler***Utsumi, I.; Mori, Y.;*

Acoustics, Speech, and Signal Processing, 1988. ICASSP-88., 1988 International Conference on , 11-14 April 1988

Page(s): 2017 -2020 vol.4

[\[Abstract\]](#) [\[PDF Full-Text \(284 KB\)\]](#) **IEEE CNF**

5 Instruction level power profiling

Mehta, H.; Owens, R.M.; Irwin, M.J.;

Acoustics, Speech, and Signal Processing, 1996. ICASSP-96. Conference Proceedings., 1996 IEEE International Conference on , Volume: 6 , 7-10 May 1996

Page(s): 3326 -3329 vol. 6

[\[Abstract\]](#) [\[PDF Full-Text \(348 KB\)\]](#) **IEEE CNF**

6 APU: specification and design of a multi algorithm ATM policing unit IP

de Lima, J.A.G.; Cavalcanti, A.C.; Ferreira de Lucena, S.;

Integrated Circuits and Systems Design, 2002. Proceedings. 15th Symposium on , 9-14 Sept. 2002

Page(s): 35 -39

[\[Abstract\]](#) [\[PDF Full-Text \(1245 KB\)\]](#) **IEEE CNF**

7 A CCITT standard 32 kbit/s ADPCM LSI codec

Nishitani, T.; Kuroda, I.; Satoh, M.; Katoh, T.; Aoki, Y.;

Acoustics, Speech, and Signal Processing [see also IEEE Transactions on Signal Processing], IEEE Transactions on , Volume: 35 Issue: 2 , Feb 1987

Page(s): 219 -225

[\[Abstract\]](#) [\[PDF Full-Text \(928 KB\)\]](#) **IEEE JNL**

8 A 300K transistor NMOS peripheral processor

Pomper, M.; Stocklinger, J.; Augspurger, U.; Mueller, B.; Horninger, K.;

Solid-State Circuits, IEEE Journal of , Volume: 19 Issue: 3 , Jun 1984

Page(s): 329 -337

[\[Abstract\]](#) [\[PDF Full-Text \(1352 KB\)\]](#) **IEEE JNL**

Terms used **microinstruction** and **power control** and **cache**

Found 936 of 113,497

Sort results by

 relevance Save results to a Binder

Display results

 expanded form Search Tips Open results in a new window[Try an Advanced Search](#)[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 Formal verification in hardware design: a survey

Christoph Kern, Mark R. Greenstreet

April 1999 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**,
Volume 4 Issue 2Full text available: [pdf\(411.53 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years, formal methods have emerged as an alternative approach to ensuring the quality and correctness of hardware designs, overcoming some of the limitations of traditional validation techniques such as simulation and testing. There are two main aspects to the application of formal methods in a design process: the formal framework used to specify desired properties of a design and the verification techniques and tools used to reason about the relationship between a spec ...

Keywords: case studies, formal methods, formal verification, hardware verification, language containment, model checking, survey, theorem proving

2 Best poster papers from MobiHoc 2002: An on-demand minimum energy routing protocol for a wireless ad hoc network

Sheetalkumar Doshi, Shweta Bhandare, Timothy X Brown

June 2002 **ACM SIGMOBILE Mobile Computing and Communications Review**, Volume 6 Issue 3Full text available: [pdf\(203.93 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

A minimum energy routing protocol reduces the energy consumption of the nodes in a wireless ad hoc network by routing packets on routes that consume the minimum amount of energy to get the packets to their destination. This paper identifies the necessary features of an *on-demand* minimum energy routing protocol and suggests mechanisms for their implementation. We highlight the importance of efficient caching techniques to store the minimum energy route information and propose the use of an ...

3 Measuring VAX 8800 performance with a histogram hardware monitor

D. W. Clark, P. J. Bannon, J. B. Keller

May 1988 **ACM SIGARCH Computer Architecture News, Proceedings of the 15th Annual International Symposium on Computer architecture**, Volume 16 Issue 2Full text available: [pdf\(1.04 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper reports the results of a study of VAX 8800 processor performance using a hardware monitor that collects histograms of the processor's micro-PC and memory bus status. The monitor keeps a count of all machine cycles executed at each micro-PC location,

as well as counting all occurrences of each bus transaction. It can measure a running system without interfering with it, and this paper's results are based on measurements of live timesharing. Because the 8800 is a microcoded machine ...

4 A processor for a high-performance personal computer

Butler W. Lampson, Kenneth A. Pier

May 1980 **Proceedings of the 7th annual symposium on Computer Architecture**

Full text available:  [pdf\(1.24 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the design goals, micro- architecture, and implementation of the microprogrammed processor for a compact high performance personal computer. This computer supports a range of high level language environments and high bandwidth I/O devices. Besides the processor, it has a cache, a memory map, main storage, and an instruction fetch unit; these are described in other papers. The processor can be shared among 16 microcoded tasks, performing microcode context switches ...

5 Performance and architectural evaluation of the PSI machine

Kazuo Taki, Katzuto Nakajima, Hiroshi Nakashima, Morihiko Ikeda

October 1987 **Proceedings of the second international conference on Architectual support for programming languages and operating systems**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Full text available:  [pdf\(1.01 MB\)](#)

Additional Information: [full citation](#), [index terms](#)

6 A Characterization of Processor Performance in the vax-11/780

Joel S. Emer, Douglas W. Clark

January 1984 **ACM SIGARCH Computer Architecture News , Proceedings of the 11th annual international symposium on Computer architecture**, Volume 12 Issue 3

Full text available:  [pdf\(980.44 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper reports the results of a study of VAX-11/780 processor performance using a novel hardware monitoring technique. A micro-PC histogram monitor was built for these measurements. It keeps a count of the number of microcode cycles executed at each microcode location. Measurement experiments were performed on live timesharing workloads as well as on synthetic workloads of several types. The histogram counts allow the calculation of the frequency of various architectural events, such as ...

7 A processor for a high-performance personal computer

Butler W. Lampson, Kenneth A. Pier

August 1998 **25 years of the international symposia on Computer architecture (selected papers)**

Full text available:  [pdf\(1.57 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

8 Multiprocessor cache analysis using ATUM

R. L. Sites, A. Agarwal

May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture**, Volume 16 Issue 2

Full text available:  [pdf\(1.38 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The design of high-performance multiprocessor systems necessitates a careful analysis of the memory system performance of parallel programs. Lacking multiprocessor address traces, previous multiprocessor performance studies using analytical models had to make an inordinate number of assumptions about the underlying memory reference patterns. We previously developed a scheme called ATUM - Address Tracing Using Microcode - to get

reliable operating system and multiprogramming traces on single ...

9 A characterization of processor performance in the VAX-11/780

Joel S. Emer, Douglas W. Clark

August 1998 **25 years of the international symposia on Computer architecture (selected papers)**

Full text available:  [pdf\(1.12 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

10 Run-time generation of HPS microinstructions from a VAX instruction stream

Y. N. Patt, S. W. Melvin, W. M. Hwu, M. C. Shebanow, C. Chen

December 1986 **ACM SIGMICRO Newsletter, Proceedings of the 19th annual workshop on Microprogramming**, Volume 17 Issue 4

Full text available:  [pdf\(808.93 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The VAX architecture is a popular ISP architecture that has been implemented in several different technologies targeted to a wide range of performance specifications. However, it has been argued that the VAX has specific characteristics which preclude a very high performance implementation. We have developed a microarchitecture (HPS) which is specifically intended for implementing very high performance computing engines. Our model of execution is a restriction on fine granularity data flow. ...

11 Resource management: Energy-efficient caching strategies in ad hoc wireless networks

Pavan Nuggehalli, Vikram Srinivasan, Carla-Fabiana Chiasserini

June 2003 **Proceedings of the fourth ACM international symposium on Mobile ad hoc networking & computing**

Full text available:  [pdf\(244.72 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we address the problem of energy-conscious cache placement in wireless ad hoc networks. We consider a network comprising a server with an interface to the wired network, and some nodes requiring access to the information stored at the server. In order to reduce access latency in such a communication environment, an effective strategy is caching the server information at some nodes distributed across the network. Caching, however, can considerably impact the system energy expenditure ...

Keywords: ad hoc networks, caching

12 Achieving high instruction cache performance with an optimizing compiler

W. W. Hwu, P. P. Chang

April 1989 **ACM SIGARCH Computer Architecture News, Proceedings of the 16th annual international symposium on Computer architecture**, Volume 17 Issue 3

Full text available:  [pdf\(1.03 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Increasing the execution power requires a high instruction issue bandwidth, and decreasing instruction encoding and applying some code improving techniques cause code expansion. Therefore, the instruction memory hierarchy performance has become an important factor of the system performance. An instruction placement algorithm has been implemented in the IMPACT-I (Illinois Microarchitecture Project using Advanced Compiler Technology - Stage I) C compiler to maximize the sequential and spatial ...

13 Cache Performance in the VAX-11/780

Douglas W. Clark

February 1983 **ACM Transactions on Computer Systems (TOCS)**, Volume 1 Issue 1

Full text available:  [pdf\(880.15 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: hardware monitor, hit ratio

14 Exploiting instruction level parallelism in processors by caching scheduled groups

Ravi Nair, Martin E. Hopkins

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture**, Volume 25 Issue 2

Full text available:  pdf(2.01 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Modern processors employ a large amount of hardware to dynamically detect parallelism in single-threaded programs and maintain the sequential semantics implied by these programs. The complexity of some of this hardware diminishes the gains due to parallelism because of longer clock period or increased pipeline latency of the machine. In this paper we propose a processor implementation which dynamically schedules groups of instructions while executing them on a fast simple engine and caches them f ...

15 Cache behavior of combinator graph reduction

Philip J. Koopman, Peter Lee, Daniel P. Siewiorek

April 1992 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 14 Issue 2

Full text available:  pdf(2.18 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The results of cache-simulation experiments with an abstract machine for reducing combinator graphs are presented. The abstract machine, called TIGRE, exhibits reduction rates that, for similar kinds of combinator graphs on similar kinds of hardware, compare favorably with previously reported techniques. Furthermore, TIGRE maps easily and efficiently onto standard computer architectures, particularly those that allow a restricted form of self-modifying code. This provides some indication th ...

Keywords: abstract machine, combinators, graph reduction, self-modifying code

16 Reducing the frequency of tag compares for low power l-cache design

Ramesh Panwar, David Rennels

April 1995 **Proceedings of the 1995 international symposium on Low power design**

Full text available:  pdf(447.34 KB) Additional Information: [full citation](#), [index terms](#)

17 Swamp: a fast processor for Smalltalk-80

David M. Lewis, David R. Galloway, Robert J. Francis, Brian W. Thomson

June 1986 **Conference proceedings on Object-oriented programming systems, languages and applications**

Full text available:  pdf(849.10 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A processor for the Smalltalk-80† programming language is described. This machine is implemented using a standard bit slice ALU and sequencer, TTL MSI, and NMOS LSI RAMS. It executes an instruction set similar to the Smalltalk-80 virtual machine instruction set. The data paths of the machine are optimized for rapid Smalltalk-80 execution by the inclusion of a context cache, tag checking, and a hardware method cache. Each context is only partly initialized when created, and has no memor ...

18 Performance from architecture: comparing a RISC and a CISC with similar hardware organization

Dileep Bhandarkar, Douglas W. Clark

19 Pipelining and performance in the VAX 8800 processor 

Douglas W. Clark

October 1987 **Proceedings of the second international conference on Architectural support for programming languages and operating systems**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

20 Compiler optimization on VLIW instruction scheduling for low power 

Chingren Lee, Jenq Kuen Lee, Tingting Hwang, Shi-Chun Tsai

April 2003 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 8 Issue 2

In this article, we investigate compiler transformation techniques regarding the problem of scheduling VLIW instructions aimed at reducing power consumption of VLIW architectures in the instruction bus. The problem can be categorized into two types: horizontal scheduling and vertical scheduling. For the case of horizontal scheduling, we propose a bipartite-matching scheme for instruction scheduling. We prove that our greedy bipartite-matching scheme always gives the optimal switching activities ...

Keywords: Compilers, VLIW instruction scheduling, instruction bus optimizations, low-power optimization

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



Terms used **microinstruction** and **power control** and **cache**

Found 936 of 113,497

Sort results
by

relevance

Save results to a Binder

Display
results

expanded form

Search Tips

Open results in a new
window

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 21 - 40 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

21 On-line architecture tuning using microcapture

A. M. Abd-Alla, Laird H. Moffett

January 1976 **ACM SIGARCH Computer Architecture News , Proceedings of the 3rd annual symposium on Computer architecture**, Volume 4 Issue 4

Full text available: [pdf\(549.12 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The modification or tuning of the micro-code in a computer that utilizes a writable control store is one method whereby a program's execution time can be improved. A method for automatically performing a microcode tuning or synthesis has been developed by Drs. Karlgaard and Abd-Alla and is discussed in detail in [1]. Presented is an extension of this effort which allows the microcode tuning to be performed on-line on program loops. This is accomplished by gathering data on characteristics ...

22 MIDL - a microinstruction description language

Marleen Sint

December 1981 **Proceedings of the 14th annual workshop on Microprogramming**

Full text available: [pdf\(848.63 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A microinstruction description language called MIDL is introduced. A MIDL description of a microarchitecture defines the semantics and triggering conditions of all microoperations. It also defines operand selection. MIDL incorporates a timing model that allows detailed specification of the timing of each microoperation, and a sequencing model that allows the description of many different sequencing schemes.

23 A message passing coprocessor for distributed memory multicomputers

Jiun-Ming Hsu, Prithviraj Banerjee

November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**

Full text available: [pdf\(1.25 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

This paper presents the architecture, methodology and performance evaluation of a *message passing coprocessor* (MPC) which can accelerate message communication in a distributed memory multicomputer. The MPC is a microprogrammable processor which off-loads the CPU of the burden of communication and speeds up the software processing by directly executing message passing instructions in microcode. It supports process scheduling, message buffer management, and fast buffer copying. The most uni ...

24 A functional level simulation engine of MAN-YO: a special purpose parallel machine for logic design automation

T. Nakata, N. Koike

June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th**

Full text available: [pdf\(511.58 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The architecture of a proto-type functional level simulator element of a massively parallel machine (MAN-YO) designed for logic design automation is presented. At functional level, hardware systems are described in a hardware description language, FDL. The FDL description is compiled into stack oriented intermediate language instructions. Communicating with other gate level/block level/ functional level processors, each functional simulator interprets the compiled instructions and simulates ...

25 A DISE implementation of dynamic code decompression

Marc L. Corliss, E. Christopher Lewis, Amir Roth

June 2003 **ACM SIGPLAN Notices , Proceedings of the 2003 ACM SIGPLAN conference on Language, compiler, and tool for embedded systems**, Volume 38 Issue 7

Full text available: [pdf\(291.52 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Code compression coupled with dynamic decompression is an important technique for both embedded and general-purpose microprocessors. *Post-fetch decompression*, in which decompression is performed after the compressed instructions have been fetched, allows the instruction cache to store compressed code but requires a highly efficient decompression implementation. We propose implementing post-fetch decompression using *dynamic instruction stream editing* (DISE), a programmable decoder-- ...

Keywords: DISE, code compression, code decompression

26 Exploiting parallel microprocessor microarchitectures with a compiler code generator

W. W. Hwu, P. P. Chang

May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture**, Volume 16 Issue 2

Full text available: [pdf\(890.51 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

With advances in VLSI technology, microprocessor designers can provide more microarchitectural parallelism to increase performance. We have identified four major forms of such parallelism: multiple microoperations issued per cycle, multiple result distribution buses, multiple execution units, and pipelined execution units. The experiments reported in this paper address two important issues: The effects of these forms and the appropriate balance among them. A central microar ...

27 Reduced instruction set computers

David A. Patterson

January 1985 **Communications of the ACM**, Volume 28 Issue 1

Full text available: [pdf\(4.06 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Reduced instruction set computers aim for both simplicity in hardware and synergy between architectures and compilers. Optimizing compilers are used to compile programming languages down to instructions that are as unencumbered as microinstructions in a large virtual address space, and to make the instruction cycle time as fast as possible.

28 An architecture for the direct execution of the Forth programming language

John R. Hayes, Martin E. Fraeman, Robert L. Williams, Thomas Zaremba

October 1987 **Proceedings of the second international conference on Architectural support for programming languages and operating systems**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Full text available: [pdf\(1.55 MB\)](#) Additional Information: [full citation](#), [citations](#), [index terms](#)

29 Using the Alfa-1 simulated processor for educational purposes

Gabriel A. Wainer, Sergio Daicz, Luis F. De Simoni, Demian Wassermann

December 2001 **Journal on Educational Resources in Computing (JERIC)**, Volume 1 Issue 4

Full text available:  [pdf\(238.65 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Alfa-1 is a simulated computer designed for computer organization courses. Alfa-1 and its accompanying toolkit allow students to acquire practical insights into developing hardware by extending existing components. The DEVS formalism is used to model individual components and to integrate them into a hierarchy that describes the detailed behavior of different levels of a computer's architecture. We introduce Alfa-1 and the toolkit, show how to extend existing components, and describe how ...

Keywords: DEVS formalism, modeling computer architectures, systems specification

30 The micro-architecture of the ECLIPSE® MV/8000: Conception and implementation

Jonathan S. Blau, Charles J. Holland, David L. Keating

November 1980 **Proceedings of the 13th annual workshop on Microprogramming**

Full text available:  [pdf\(622.95 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The microcode of the ECLIPSE MV/8000 controls the hardware to emulate an instruction set. In the MV/8000 the micro-architecture is defined and limited by the following constraints: 1) the desire to implement microcode in a limited number of locations; 2) the use of LSI technology; 3) a virtual memory architecture. This paper will attempt to show how each of these factors contributed to the micro-architecture, to describe that architecture, and to relate ...

31 Evaluation of a high performance code compression method

Charles Lefurgy, Eva Piccininni, Trevor Mudge

November 1999 **Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture**

Full text available:   [pdf\(1.01 MB\)](#)  Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)
[Publisher Site](#)

Compressing the instructions of an embedded program is important for cost-sensitive low-power control-oriented embedded computing. A number of compression schemes have been proposed to reduce program size. However, the increased instruction density has an accompanying performance cost because the instructions must be decompressed before execution. In this paper, we investigate the performance penalty of a hardware-managed code compression algorithm recently introduced in IBM's PowerPC 405. ...

32 Performance effects of architectural complexity in the Intel 432

Robert P. Colwell, Edward F. Gehringer, E. Douglas Jensen

August 1988 **ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 3

Full text available:  [pdf\(3.45 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Intel 432 is noteworthy as an architecture incorporating a large amount of functionality that most other systems perform by software. It has, in effect, "migrated" this functionality from the software into the microcode and hardware. The benefits of functional migration have recently been a subject of intense controversy, with critics claiming that a complex architecture is inherently less efficient than a simple architecture with good software support. This paper examines the ...

33 A fill-unit approach to multiple instruction issue

Manoj Franklin, Mark Smotherman

November 1994 **Proceedings of the 27th annual international symposium on Microarchitecture**

Additional Information:

Full text available: [pdf\(992.94 KB\)](#)

[full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Multiple issue of instructions occurs in superscalar and VLIW machines. This paper investigates a third type of machine design, which combines the advantages of code compatibility as in superscalars and the absence of complex dependency-checking logic from the decoder as in VLIW. In this design, a stream of scalar instructions is executed by the hardware and is simultaneously compacted into VLIW-type instructions, which are then stored in a structure called a shadow cache. When a shadow cache ...

Keywords: VLIW, instruction-level parallelism, multiple operation issue, superscalar

34 Design considerations for the VLSI processor of X-TREE

David A. Patterson, E. Scott Fehr, Carlo H. Séquin

April 1979 **Proceedings of the 6th annual symposium on Computer architecture**

Full text available: [pdf\(958.46 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

X-NODE is a single-chip VLSI processor to be realized in the mid 1980's and to be used as a building block for a tree-structured multiprocessor system (X-TREE). Three major trends influence the design of this processor: the continuing evolution of VLSI technology, the requirements for parallelism and communication in a multiprocessor system, and the need for better support of software and high level language constructs. The influence of these trends on the processor architecture are discussed ...

35 Multiple instruction issue in the NonStop cyclone processor

Robert W. Horst, Richard L. Harris, Robert L. Jardine

May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture**, Volume 18 Issue 3

Full text available: [pdf\(1.06 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the architecture for issuing multiple instructions per clock in the NonStop Cyclone Processor. Pairs of instructions are fetched and decoded by a dual two-stage prefetch pipeline and passed to a dual six-stage pipeline for execution. Dynamic branch prediction is used to reduce branch penalties. A unique microcode routine for each pair is stored in the large duplexed control store. The microcode controls parallel data paths optimized for executing the most frequent instructions ...

36 Difficult-path branch prediction using subordinate microthreads

Robert S. Chappell, Francis Tseng, Adi Yoaz, Yale N. Patt

May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture**, Volume 30 Issue 2

Full text available: [pdf\(1.14 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

Branch misprediction penalties continue to increase as microprocessor cores become wider and deeper. Thus, improving branch prediction accuracy remains an important challenge. Simultaneous Subordinate Microthreading (SSMT) provides a means to improve branch prediction accuracy. SSMT machines run multiple, concurrent microthreads in support of the primary thread. We propose to dynamically construct microthreads that can speculatively and accurately pre-compute branch outcomes along frequently mis ...

37 ByteLisp and its Alto implementation

L. Peter Deutsch

August 1980 **Proceedings of the 1980 ACM conference on LISP and functional programming**

Full text available: [pdf\(989.19 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes in detail the most interesting aspects of ByteLisp, a transportable Lisp system architecture which implements the Interlisp dialect of Lisp, and its first

implementation, on a microprogrammed minicomputer called the Alto. Two forthcoming related papers will deal with general questions of Lisp machine and system architecture, and detailed measurements of the Alto ByteLisp system described here. A highly condensed summary of the series was published at MICRO-11 in November ...

38 PSCP: a scalable parallel ASIP architecture for reactive systems

A. Pyttel, A. Sedlmeier, C. Veith

February 1998 **Proceedings of the conference on Design, automation and test in Europe**

Full text available:  [pdf\(208.55 KB\)](#)

 [Publisher Site](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We describe a Codesign approach based on a parallel and scalable ASIP architecture, which is suitable for the implementation of reactive systems. The specification language of our approach is extended statecharts. Our ASIP architecture is scalable with respect to the number of processing elements as well as parameters such as bus widths and register file sizes. Instruction sets are generated from a library of components covering a spectrum of space/time trade-off alternatives. Our approach featu ...

Keywords: FPGA, application-specific, statechart, modular

39 Virtual memory on a narrow machine for an object-oriented language

Ted Kaehler

June 1986 **Conference proceedings on Object-oriented programming systems, languages and applications**

Full text available:  [pdf\(1.66 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

LOOM (Large Object-Oriented Memory) is a virtual memory implemented in software that supports the Smalltalk-80™ programming language and environment on the Xerox Dorado computer. LOOM provides 8 billion bytes of secondary memory address space and is specifically designed to run on computers with a narrow word size (16-bit wide words). All storage is viewed as objects that contain fields. Objects may have an average size as small as 10 fields. LOOM swaps objects between primary and s ...

40 Branch folding in the CRISP microprocessor: reducing branch delay to zero

D. R. Ditzel, H. R. McLellan

June 1987 **Proceedings of the 14th annual international symposium on Computer architecture**

Full text available:  [pdf\(705.60 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A new method of implementing branch instructions is presented. This technique has been implemented in the CRISP Microprocessor. With a combination of hardware and software techniques the execution time cost for many branches can be effectively reduced to zero. Branches are folded into other instructions, making their execution as separate instructions unnecessary. Branch Folding can reduce the apparent number of instructions needed to execute a program by the number of bran ...

Results 21 - 40 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)

 [QuickTime](#)

 [Windows Media Player](#)

 [Real Player](#)

Terms used **microinstruction** and **power control** and **cache**

Found 936 of 113,497

Sort results
by relevance Save results to a BinderTry an [Advanced Search](#)Display
results expanded form Search TipsTry this search in [The ACM Guide](#) Open results in a new window

Results 41 - 60 of 200

Result page: [previous](#)[1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

41 An environment for research in microprogramming and emulation

Robert F. Rosin, Gideon Frieder, Richard H. Eckhouse

August 1972 **Communications of the ACM**, Volume 15 Issue 8Full text available: [pdf\(1.33 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The development of the research project in microprogramming and emulation at State University of New York at Buffalo consisted of three phases: the evaluation of various possible machines to support this research; the decision to purchase one such machine, which appears to be superior to the others considered; and the organization and definition of goals for each group in the project. Each of these phases is reported, with emphasis placed on the early results achieved in this research.

Keywords: computer systems, emulation, hardware evaluation, input-output systems, language processors, microprogramming, nanoprogram, project management

42 Phase coupling and constant generation in an optimizing microcode compiler

Steven R. Vegdahl

October 1982 **Proceedings of the 15th annual workshop on Microprogramming**Full text available: [pdf\(831.84 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The designer of an optimizing compiler must concern himself with the order in which optimization phases are performed; a pair of phases may be interdependent in the sense that each phase could benefit from information produced by the other. In a compiler for a horizontal target architecture, one such phase-ordering problem occurs between code-generation and compaction. Presented here is an overview of a research effort at Carnegie-Mellon University which ha ...

43 Performance of Lisp systems

Richard P. Gabriel, Larry M. Masinter

August 1982 **Proceedings of the 1982 ACM symposium on LISP and functional programming**Full text available: [pdf\(1.45 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the issues involved in evaluating the performance of Lisp systems. We explore the various levels at which quantitative statements can be made about the performance of a Lisp system, giving examples from existing implementations wherever possible. Our thesis is that benchmarking is most effective when performed in conjunction with an analysis of the underlying Lisp implementation and computer architecture. We examine some simple benchmarks which have been used to measure ...

44 [Session 3: Energy-aware OS's: The benefits of event-driven energy accounting in power-sensitive systems](#)

Frank Bellosa

September 2000 **Proceedings of the 9th workshop on ACM SIGOPS European workshop: beyond the PC: new challenges for the operating system**

Full text available:  [pdf\(86.80 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

A prerequisite of energy-aware scheduling is precise knowledge of any activity inside the computer system. Embedded hardware monitors (e.g., processor performance counters) have proved to offer valuable information in the field of performance analysis. The same approach can be applied to investigate the energy usage patterns of individual threads. We use information about active hardware units (e.g., integer/floating-point unit, cache/memory interface) gathered by event counters to establish a t ...

45 [Performance of the VAX-11/780 translation buffer: simulation and measurement](#)

Douglas W. Clark, Joel S. Emer

February 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 1

Full text available:  [pdf\(2.36 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A virtual-address translation buffer (TB) is a hardware cache of recently used virtual-to-physical address mappings. The authors present the results of a set of measurements and simulations of translation buffer performance in the VAX-11/780. Two different hardware monitors were attached to VAX-11/780 computers, and translation buffer behavior was measured. Measurements were made under normal time-sharing use and while running reproducible synthetic time-sharing work loads. Reported measure ...

46 [Code generation and scheduling: Compiler optimization on instruction scheduling for low power](#)

Chingren Lee, Jenq Kuen Lee, TingTing Hwang, Shi-Chun Tsai

September 2000 **Proceedings of the 13th international symposium on System synthesis**

Full text available:  [pdf\(375.33 KB\)](#) Additional Information: [full citation](#), [abstract](#)

In this paper, we investigate the compiler transformation techniques to the problem of scheduling VLIW instructions aimed to reduce the power consumption on the instruction bus. It can be categorized into two types: horizontal and vertical scheduling. For the horizontal case, we propose a bipartite-matching scheme. We prove that our greedy algorithm always gives the optimal switching activities of the instruction bus. In the vertical case, we prove that the problem is NP-hard, and propose a heur ...

47 [Firefly: a multiprocessor workstation](#)

Charles P. Thacker, Lawrence C. Stewart

October 1987 **Proceedings of the second international conference on Architectural support for programming languages and operating systems**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Full text available:  [pdf\(1.10 MB\)](#) Additional Information: [full citation](#), [citations](#), [index terms](#)

48 [The hardware architecture of the CRISP microprocessor](#)

D. R. Ditzel, H. R. McLellan, A. D. Berenbaum

June 1987 **Proceedings of the 14th annual international symposium on Computer architecture**

Full text available:  [pdf\(930.17 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

49 [A trace-driven simulation methodology](#)

Humayun Khalid

This paper presents a simulation methodology for evaluating the performance of CISC computers. The method is called Message Flow Technique (MFT). MFT has several advantages over Instruction Flow Technique (IFT) we presented in [1]. The proposed methodology is applied to a single and two-level cache CISC system using 80486 SX as a case study. It was found that with a single-level on-chip cache of size 8K, the performance of the system is considerably limited by the service time of BIU(Bus In ...

Keywords: gibson mix, modified, multi-level cache, performance evaluation, simulation methodologies

50 [ATUM: a new technique for capturing address traces using microcode](#)

A. Agarwal, R. L. Sites, M. Horowitz

June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture**, Volume 14 Issue 2

Full text available: [pdf\(894.10 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Trace-driven simulation is often used in the design of computer systems, especially caches and translation lookaside buffers. Capturing address traces to drive such simulations has been problematic, often involving 1000:1 software overhead to trace a target workload, and/or mechanisms that cause significant distortions in the recorded data. A new technique for capturing address traces has been developed to use a processor's microcode to record addresses in a reserved part of main memory as ...

51 [Executing compressed programs on an embedded RISC architecture](#)

Andrew Wolfe, Alex Chanin

December 1992 **ACM SIGMICRO Newsletter , Proceedings of the 25th annual international symposium on Microarchitecture**, Volume 23 Issue 1-2

Full text available: [pdf\(1.53 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

52 [Massive arrays of idle disks for storage archives](#)

Dennis Colarelli, Dirk Grunwald

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Full text available: [pdf\(751.87 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

The declining costs of commodity disk drives is rapidly changing the economics of deploying large amounts of online or near-line storage. Conventional mass storage systems use either high performance RAID clusters, automated tape libraries or a combination of tape and disk. In this paper, we analyze an alternative design using *massive arrays of idle disks*, or MAID. We argue that this storage organization provides storage densities matching or exceeding those of tape libraries with perform ...

53 [Design tradeoffs to support the C programming language in the CRISP microprocessor](#)

David R. Ditzel, Hubert R. McLellan, Alan D. Berenbaum

October 1987 **Proceedings of the second international conference on Architectural support for programming languages and operating systems**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Full text available: [pdf\(688.92 KB\)](#) Additional Information: [full citation](#), [citations](#), [index terms](#)

54 [The architecture of the hardware unification unit and an implementation](#)

N. S. Woo

December 1985 **ACM SIGMICRO Newsletter , Proceedings of the 18th annual workshop**

This paper describes the architecture and the current implementation of the hardware unification unit (HUU). The HUU performs the literal unification operation in Prolog processing. It is designed as a coprocessor to a host system that handles other operations of Prolog processing such as bookkeeping and sequencing. After the host system provides input values to the HUU and activates it, the HUU works independently from the host system; when it finishes its operation it reports the result to ...

55 Mobility: PATHS: analysis of PATH duration statistics and their impact on reactive MANET routing protocols

Narayanan Sadagopan, Fan Bai, Bhaskar Krishnamachari, Ahmed Helmy

June 2003 **Proceedings of the fourth ACM international symposium on Mobile ad hoc networking & computing**

Full text available:  pdf(311.57 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We develop a detailed approach to study how mobility impacts the performance of reactive MANET routing protocols. In particular we examine how the statistics of path durations including PDFs vary with the parameters such as the mobility model, relative speed, number of hops, and radio range. We find that at low speeds, certain mobility models may induce multi-modal distributions that reflect the characteristics of the spatial map, mobility constraints and the communicating traffic pattern. However ...

Keywords: mobile ad hoc network, mobility, path duration, performance

56 The Clipper processor: instruction set architecture and implementation

W. Hollingsworth, H. Sachs, A. J. Smith

February 1989 **Communications of the ACM**, Volume 32 Issue 2

Full text available:  pdf(4.67 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Intergraph's CLIPPER microprocessor is a high performance, three chip module that implements a new instruction set architecture designed for convenient programmability, broad functionality, and easy future expansion.

57 Firmware approach to fast Lisp interpreter

Hiroshi G. Okuno, Nobuyasu Osato, Ikuo Takeuchi

December 1987 **Proceedings of the 20th annual workshop on Microprogramming**

Full text available:  pdf(1.14 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The approach to speed up a Lisp interpreter by implementing it in firmware seems promising. A microcoded Lisp interpreter shows good performance for very simple benchmarks, while it often fails to provide good performance for larger benchmarks and applications unless speedup techniques are devised for it. This was the case for the TAO/ELIS system. This paper describes various techniques devised for the TAO/ELIS system in order to speed up the interpreter of the TAO language implemented on to ...

58 Towards an efficient, machine-independent language for microprogramming

David A. Patterson, Karl Lew, Richard Tuck

November 1979 **Proceedings of the 12th annual workshop on Microprogramming**

Full text available:  pdf(913.17 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A machine independent low level language YALLL is presented. This language produces microcode for two very different machines: Hewlett Packard HP 300 and Digital Equipment Corporation VAX 11/780. The efficiency of this language is tested by comparing two examples on both machines to microassembly coded versions. To our best knowledge, this

is the first time programs have been compiled and executed on two different microarchitectures. These examples also let us compare the efficiency of the ...

59 On tuning the microarchitecture of an HPS implementation of the VAX

James E. Wilson, Steve Melvin, Michael Shebanow, Wen-mei Hwu, Yale N. Patt

December 1987 **Proceedings of the 20th annual workshop on Microprogramming**

Full text available:  [pdf\(744.71 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The HPS Microarchitecture has been developed as an execution model for implementing various architectures at very high performance. A considerable amount of effort has gone into the use of HPS as a microarchitecture for the VAX. In this paper, we describe our first full simulation of the microVAX subset, and report the results of varying (i.e. tuning) certain important parameters.

60 Self-assessment procedure XVI: a self-assessment procedure dealing with computer

organization and logic design

Glen G. Langdon

November 1986 **Communications of the ACM**, Volume 29 Issue 11

Full text available:  [pdf\(908.13 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A self-assessment procedure dealing with computer organization and logic design

Results 41 - 60 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



Terms used **microinstruction** and **power control** and **cache** and **disable** or **disabled** or **disabling**

Found 4,365 of 113,497

Sort results
by

relevance

Save results to a Binder

[Try an Advanced Search](#)

Display
results

expanded form

Search Tips

[Try this search in The ACM Guide](#)

Open results in a new
window

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [Cache Performance in the VAX-11/780](#)

Douglas W. Clark

February 1983 **ACM Transactions on Computer Systems (TOCS)**, Volume 1 Issue 1

Full text available: [pdf\(880.15 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



Keywords: hardware monitor, hit ratio

2 [Best poster papers from MobiHoc 2002: An on-demand minimum energy routing protocol for a wireless ad hoc network](#)

Sheetalkumar Doshi, Shweta Bhandare, Timothy X Brown

June 2002 **ACM SIGMOBILE Mobile Computing and Communications Review**, Volume 6 Issue 3

Full text available: [pdf\(203.93 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)



A minimum energy routing protocol reduces the energy consumption of the nodes in a wireless ad hoc network by routing packets on routes that consume the minimum amount of energy to get the packets to their destination. This paper identifies the necessary features of an *on-demand* minimum energy routing protocol and suggests mechanisms for their implementation. We highlight the importance of efficient caching techniques to store the minimum energy route information and propose the use of an ...

3 [Selective cache ways: on-demand cache resource allocation](#)

David H. Albonesi

November 1999 **Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture**

Full text available: [pdf\(1.09 MB\)](#) [Publisher Site](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



Increasing levels of microprocessor power dissipation call for new approaches at the architectural level that save energy by better matching of on-chip resources to application requirements. Selective cache ways provides the ability to disable a subset of the ways in a set associative cache during periods of modest cache activity, while the full cache may remain operational for more cache-intensive periods. Because this approach leverages the subarray partitioning that is already pr ...

4 [Multiprocessor cache analysis using ATUM](#)

R. L. Sites, A. Agarwal



The design of high-performance multiprocessor systems necessitates a careful analysis of the memory system performance of parallel programs. Lacking multiprocessor address traces, previous multiprocessor performance studies using analytical models had to make an inordinate number of assumptions about the underlying memory reference patterns. We previously developed a scheme called ATUM - Address Tracing Using Microcode - to get reliable operating system and multiprogramming traces on single ...

5 **Reducing the frequency of tag compares for low power l-cache design**

Ramesh Panwar, David Rennels

April 1995 **Proceedings of the 1995 international symposium on Low power design**

Full text available: [pdf\(447.34 KB\)](#) Additional Information: [full citation](#), [index terms](#)



6 **The white dwarf: a high-performance application-specific processor**

A. Wolfe, M. Breternitz, C. Stephens, A. L. Ting, D. B. Kirk, R. P. Bianchini, J. P. Shen

May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture**, Volume 16 Issue 2

Full text available: [pdf\(1.40 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents the design and implementation of a high-performance special-purpose processor, called The White Dwarf, for accelerating finite element analysis algorithms. The White Dwarf CPU contains two Am29325 32-bit floating-point processors and one Am29332 32-bit ALU, and employs a wide-instruction word architecture in which the application algorithm is directly implemented in microcode. The entire system is VME-bus compatible and interfaces with a SUN 31160 host. The syste ...



7 **A message passing coprocessor for distributed memory multicomputers**

Jiun-Ming Hsu, Prithviraj Banerjee

November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**

Full text available: [pdf\(1.25 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

This paper presents the architecture, methodology and performance evaluation of a *message passing coprocessor* (MPC) which can accelerate message communication in a distributed memory multicomputer. The MPC is a microprogrammable processor which off-loads the CPU of the burden of communication and speeds up the software processing by directly executing message passing instructions in microcode. It supports process scheduling, message buffer management, and fast buffer copying. The most uni ...



8 **Performance of the VAX-11/780 translation buffer: simulation and measurement**

Douglas W. Clark, Joel S. Emer

February 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 1

Full text available: [pdf\(2.36 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A virtual-address translation buffer (TB) is a hardware cache of recently used virtual-to-physical address mappings. The authors present the results of a set of measurements and simulations of translation buffer performance in the VAX-11/780. Two different hardware monitors were attached to VAX-11/780 computers, and translation buffer behavior was measured. Measurements were made under normal time-sharing use and while running reproducible synthetic time-sharing work loads. Reported measure ...



9 **The Clipper processor: instruction set architecture and implementation**

W. Hollingsworth, H. Sachs, A. J. Smith

Intergraph's CLIPPER microprocessor is a high performance, three chip module that implements a new instruction set architecture designed for convenient programmability, broad functionality, and easy future expansion.

10 A DISE implementation of dynamic code decompression

Marc L. Corliss, E. Christopher Lewis, Amir Roth

June 2003 **ACM SIGPLAN Notices , Proceedings of the 2003 ACM SIGPLAN conference on Language, compiler, and tool for embedded systems**, Volume 38 Issue 7

Full text available:  [pdf\(291.52 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Code compression coupled with dynamic decompression is an important technique for both embedded and general-purpose microprocessors. *Post-fetch decompression*, in which decompression is performed after the compressed instructions have been fetched, allows the instruction cache to store compressed code but requires a highly efficient decompression implementation. We propose implementing post-fetch decompression using *dynamic instruction stream editing* (DISE), a programmable decoder-- ...

Keywords: DISE, code compression, code decompression

11 Predicting the usefulness of a block result: a micro-architectural technique for high-performance low-power processors

Enric Musoll

November 1999 **Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture**

Full text available:   [pdf\(1.09 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper proposes a micro-architectural technique in which a prediction is made for some power-hungry units of a processor. The prediction consists of whether the result of a particular unit or block of logic will be useful in order to execute the current instruction. If it is predicted useless, then that block is disabled. It would be ideal if the predictions were totally accurate, thus not decreasing the instruction-per-cycle (IPC) performance metric. ...

12 Testing and Debugging Custom Integrated Circuits

Edward H. Frank, Robert F. Sproull

December 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 4

Full text available:  [pdf\(2.25 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

13 Using the Alfa-1 simulated processor for educational purposes

Gabriel A. Wainer, Sergio Daicz, Luis F. De Simoni, Demian Wassermann

December 2001 **Journal on Educational Resources in Computing (JERIC)**, Volume 1 Issue 4

Full text available:  [pdf\(238.65 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Alfa-1 is a simulated computer designed for computer organization courses. Alfa-1 and its accompanying toolkit allow students to acquire practical insights into developing hardware by extending existing components. The DEVS formalism is used to model individual components and to integrate them into a hierarchy that describes the detailed behavior of different levels of a computer's architecture. We introduce Alfa-1 and the toolkit, show how to extend existing components, and describe how ...

Keywords: DEVS formalism, modeling computer architectures, systems specification

14 Performance effects of architectural complexity in the Intel 432

Robert P. Colwell, Edward F. Gehringer, E. Douglas Jensen

August 1988 **ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 3

Full text available:  [pdf\(3.45 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Intel 432 is noteworthy as an architecture incorporating a large amount of functionality that most other systems perform by software. It has, in effect, "migrated" this functionality from the software into the microcode and hardware. The benefits of functional migration have recently been a subject of intense controversy, with critics claiming that a complex architecture is inherently less efficient than a simple architecture with good software support. This paper examines t ...

15 Difficult-path branch prediction using subordinate microthreads

Robert S. Chappell, Francis Tseng, Adi Yoaz, Yale N. Patt

May 2002 **ACM SIGARCH Computer Architecture News, Proceedings of the 29th annual international symposium on Computer architecture**, Volume 30 Issue 2

Full text available:  [pdf\(1.14 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

Branch misprediction penalties continue to increase as microprocessor cores become wider and deeper. Thus, improving branch prediction accuracy remains an important challenge. Simultaneous Subordinate Microthreading (SSMT) provides a means to improve branch prediction accuracy. SSMT machines run multiple, concurrent microthreads in support of the primary thread. We propose to dynamically construct microthreads that can speculatively and accurately pre-compute branch outcomes along frequently mis ...

16 Performance of Lisp systems

Richard P. Gabriel, Larry M. Masinter

August 1982 **Proceedings of the 1982 ACM symposium on LISP and functional programming**

Full text available:  [pdf\(1.45 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the issues involved in evaluating the performance of Lisp systems. We explore the various levels at which quantitative statements can be made about the performance of a Lisp system, giving examples from existing implementations wherever possible. Our thesis is that benchmarking is most effective when performed in conjunction with an analysis of the underlying Lisp implementation and computer architecture. We examine some simple benchmarks which have been used to measure ...

17 A trace-driven simulation methodology

Humayun Khalid

December 1995 **ACM SIGARCH Computer Architecture News**, Volume 23 Issue 5

Full text available:  [pdf\(573.30 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

This paper presents a simulation methodology for evaluating the performance of CISC computers. The method is called Message Flow Technique (MFT). MFT has several advantages over Instruction Flow Technique (IFT) we presented in [1]. The proposed methodology is applied to a single and two-level cache CISC system using 80486 SX as a case study. It was found that with a single-level on-chip cache of size 8K, the performance of the system is considerably limited by the service time of BIU(Bus In ...

Keywords: gibson mix, modified, multi-level cache, performance evaluation, simulation methodologies

18 ATUM: a new technique for capturing address traces using microcode

A. Agarwal, R. L. Sites, M. Horowitz

Trace-driven simulation is often used in the design of computer systems, especially caches and translation lookaside buffers. Capturing address traces to drive such simulations has been problematic, often involving 1000:1 software overhead to trace a target workload, and/or mechanisms that cause significant distortions in the recorded data. A new technique for capturing address traces has been developed to use a processor's microcode to record addresses in a reserved part of main memory as ...

19 Design of a user-microprogrammable building block

Michael Kraley, Randall Rettberg, Philip Herman, Robert Bressler, Anthony Lake
November 1980 **Proceedings of the 13th annual workshop on Microprogramming**

A user-microprogrammable computer has been developed for use as a building block in general-purpose and dedicated computer systems. The architecture is designed to be easily microprogrammed and features a 32-bit, vertically oriented microinstruction. The processor has a 135-nanosecond cycle time, either 16- or 20-bit macro data paths, and 1024 hardware registers. A significant fraction of the processor bandwidth may be budgeted for I/O processing to allow the substitution of microcode for e ...

20 Architecture of SOAR: Smalltalk on a RISC

David Ungar, Ricki Blau, Peter Foley, Dain Samples, David Patterson
January 1984 **ACM SIGARCH Computer Architecture News , Proceedings of the 11th annual international symposium on Computer architecture**, Volume 12 Issue 3

Smalltalk on a RISC (SOAR) is a simple, Von Neumann computer that is designed to execute the Smalltalk-80 system much faster than existing VLSI microcomputers. The Smalltalk-80 system is a highly productive programming environment but poses tough challenges for implementors: dynamic data typing, a high level instruction set, frequent and expensive procedure calls, and object-oriented storage management. SOAR compiles programs to a low level, efficient instruction set. Parallel tag checks pe ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)



Selective cache ways: on-demand cache resource allocation

Full text [Publisher Site](#) [Pdf \(1.09 MB\)](#)**Source**

International Symposium on Microarchitecture [archive](#)
Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture
[table of contents](#)
Haifa, Israel
Pages: Pages: 248 - 259
Year of Publication: 1999
ISBN:0-7695-0437-X

Author [David H. Albonesi](#)**Sponsors** IEEE TC - MICRO : IEEE TC - MICRO

[SIGMICRO](#): ACM Special Interest Group on Microarchitectural Research and Processing

Publisher IEEE Computer Society Washington, DC, USA

Additional Information: [abstract](#) [references](#) [citations](#) [index terms](#) [peer to peer](#)

Tools and Actions:

[Discussions](#) [Find similar Articles](#) [Review this Article](#)
[Save this Article to a Binder](#) [Display in BibTeX Format](#)

↑ ABSTRACT

Increasing levels of microprocessor power dissipation call for new approaches at the architectural level that save energy by better matching of on-chip resources to application requirements. Selective cache ways provides the ability to disable a subset of the ways in a set associative cache during periods of modest cache activity, while the full cache may remain operational for more cache-intensive periods. Because this approach leverages the subarray partitioning that is already present for performance reasons, only minor changes to a conventional cache are required, and therefore, full-speed cache operation can be maintained. Furthermore, the tradeoff between performance and energy is flexible, and can be dynamically tailored to meet changing application and machine environmental conditions. We show that trading off a small performance degradation for energy savings can produce a significant reduction in cache energy dissipation using this approach.

↑ REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

1 [David H. Albonesi, Dynamic IPC/clock rate optimization, Proceedings of the 25th annual international symposium on Computer architecture, p.282-292, June 27-July 02, 1998, Barcelona, Spain](#)

2 [Glenn Ammons , Thomas Ball , James R. Larus, Exploiting hardware performance counters with flow and context sensitive profiling, Proceedings of the 1997 ACM SIGPLAN conference on Programming language design and implementation, p.85-96, June 16-18, 1997, Las Vegas, Nevada, United States](#)

3 Jennifer M. Anderson , Lance M. Berc , Jeffrey Dean , Sanjay Ghemawat , Monika R. Henzinger , Shun-Tak A. Leung , Richard L. Sites , Mark T. Vandevenne , Carl A. Waldspurger , William E. Weihl , Continuous profiling: where have all the cycles gone?, Proceedings of the sixteenth ACM symposium on Operating systems principles , p.1-14, October 05-08, 1997, Saint Malo, France

4 Anant Agarwal , John Hennessy , Mark Horowitz , Cache performance of operating system and multiprogramming workloads, ACM Transactions on Computer Systems (TOCS) , v.6 n.4, p.393-431, Nov. 1988

5 P. Bannon. Alpha 21364: A scalable single-chip SMP. Microprocessor Forum, October 1998.

6 Nikolaos Bellas Ibrahim Hajj , George Stamoulis , N. Bellas , C. Polychronopoulos , Architectural and compiler support for energy reduction in the memory hierarchy of high performance microprocessors, Proceedings 1998 international symposium on Low power electronics and design , p.70-75, August 10-12, 1998, Monterey, California, United States

7 William J. Bowhill , Shane L. Bell , Bradley J. Benschneider , Andrew J. Black , Sharon M. Britton , Ruben W. Castelino , Dale R. Donchin , John H. Edmondson , Harry R. Fair , Paul E. Gronowski , Anil K. Jain , Patricia L. Kroesen , Marc E. Lamere , Bruce J. Loughlin , Shekhar Mehata , Sribalan Santhanam , Timothy A. Shedd , Stephen C. Thierauf , Robert O. Mueller , Ronald P. Preston , Michael J. Smith , Circuit implementation of a 300-MHz 64-bit second-generation CMOS Alpha CPU , Digital Technical Journal , v.7 n.1, p.100-118, Jan. 1995

8 D. Burger and T. Austin. The SimpleScalar toolset, version 2.0. Technical Report TR-97-1342, University of Wisconsin-Madison, June 1997.

9 A. Chandrakasan, S. Sheng, and R. Brodersen. Low-power CMOS digital design. IEEE Journal of Solid-State Circuits, 27(4):473-484, April 1992.

10 Jeffrey Dean , James E. Hicks , Carl A. Waldspurger , William E. Weihl , George Chrysos , ProfileMe: hardware support for instruction-level profiling on out-of-order processors, Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture , p.292-302, December 01-03, 1997, Research Triangle Park, North Carolina, United States

11 P. Dinda et al. The CMU task parallel program suite. Technical Report CMU-CS-94-131, Carnegie Mellon University, March 1994.

12 D. Dobberpuhl et al. A 200MHz, 64-bit, dual-issue CMOS microprocessor. Digital Technical Journal, 4(4):35-50, Special Issue 1992.

13 John H. Edmondson , Paul I. Rubinfeld , Peter J. Bannon , Bradley J. Benschneider , Debra Bernstein , Ruben W. Castelino , Elizabeth M. Cooper , Daniel E. Dever , Dale R. Donchin , Timothy C. Fischer , Anil K. Jain , Shekhar Mehta , Jeanne E. Meyer , Ronald P. Preston , Vidya Rajagopalan , Chandrasekhara Somanathan , Scott A. Taylor , Gilbert M. Wolrich , Internal organization of the Alpha 21164, a 300-MHz 64-bit quad-issue CMOS RISC microprocessor , Digital Technical Journal , v.7 n.1, p.119-135, Jan. 1995

14 T. Horel and G. Lauterbach. UltraSPARC III: Designing third-generation 64-bit performance. IEEE Micro, 19(3):73- 85, May/June 1999.

15 Milind B. Kamble , Kanad Ghose , Analytical energy dissipation models for low-power caches , Proceedings of the 1997 international symposium on Low power electronics and design , p.143-148, August 18-20, 1997, Monterey, California, United States

16 R. Kessler. The Alpha 21264 microprocessor. IEEE Micro, 19(2):24-36, March/April 1999.

17 R. Kessler, E. McLellan, and D. Webb. The Alpha 21264 microprocessor architecture. International Conference on Computer Design, October 1998.

18 Johnson Kin , Munish Gupta , William H. Mangione-Smith, The filter cache: an energy efficient memory structure, Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture, p.184-193, December 01-03, 1997, Research Triangle Park, North Carolina, United States

19 A. Kumar. The HP PA-8000 RISC CPU. IEEE Computer, 17(2):27-32, March 1997.

20 A. Lebeck and D. Wood. Cache profiling and the SPEC benchmarks: A case study. IEEE Computer, 27(10):15-26, October 1994.

21 G. Lesartre and D. Hunt. PA-8500: The continuing evolution of the PA-8000 family. Proceedings of Compcon, 1997.

22 E. McLellan. The Alpha AXP architecture and 21064 processor. IEEE Micro, 13(4):36-47, June 1993.

23 James Montanaro , Richard T. Witek , Krishna Anne , Andrew J. Black , Elizabeth M. Cooper , Daniel W. Dobberpuhl , Paul M. Donahue , Jim Eno , Gregory W. Hoeppner , David Kruckemyer , Thomas H. Lee , Peter C. M. Lin , Liam Madden , Daniel Murray , Mark H. Pearce , Sribalan Santhanam , Kathryn J. Snyder , Ray Stephany , Stephen C. Thierauf, A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor, Digital Technical Journal, v.9 n.1, p.49-62, 1997

24 M. Tremblay and J. O'Connor. UltraSparc I: A four-issue processor supporting multimedia. IEEE Micro, 16(2):42- 50, April 1996.

25 S. Wilton and N. Jouppi. An enhanced access and cycle time model for on-chip caches. Technical Report 93/5, Digital Western Research Laboratory, July 1994.

26 Xiaolan Zhang , Zheng Wang , Nicholas Gloy , J. Bradley Chen , Michael D. Smith, System support for automatic profiling and optimization, Proceedings of the sixteenth ACM symposium on Operating systems principles, p.15-26, October 05-08, 1997, Saint Malo, France

↑ CITINGS 26

S. Kim , N. Vijaykrishnan , M. Kandemir , M. J. Irwin, Energy-efficient instruction cache using page-based placement, Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems, November 16-17, 2001, Atlanta, Georgia, USA

Victor Delaluz , Mahmut Kandemir , N. Vijaykrishnan , Anand Sivasubramaniam , Mary Jane Irwin, Hardware and Software Techniques for Controlling DRAM Power Modes, IEEE Transactions on Computers, v.50 n.11, p.1154-1173, November 2001

J. S. Hu , M. Kandemir , N. Vijaykrishnan , M. J. Irwin , H. Saputra , W. Zhang, Compiler-directed cache polymorphism, ACM SIGPLAN Notices, v.37 n.7, July 2002

Zhigang Hu , Stefanos Kaxiras , Margaret Martonosi, Let caches decay: reducing leakage energy via exploitation of cache generational behavior, ACM Transactions on Computer Systems (TOCS), v.20 n.2, p.161-190, May 2002

Hongbo Yang , Guang R. Gao , Clement Leung, On achieving balanced power consumption in software pipelined loops, Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems, October 08-11, 2002, Greenoble, France

S. Kim , N. Vijaykrishnan , M. Kandemir , A. Sivasubramaniam , M. J. Irwin , E. Geethanjali, Power-aware partitioned cache architectures, Proceedings of the 2001 international symposium on Low power electronics and design, p.64-67, August 2001, Huntington Beach, California, United States

Michael Powell , Se-Hyun Yang , Babak Falsafi , Kaushik Roy , T. N. Vijaykumar, Gated-Vdd: a circuit technique to reduce leakage in deep-submicron cache memories, Proceedings of the 2000 international symposium on Low power electronics and design, p.90-95, July 25-27, 2000, Rapallo, Italy

Michael Zhang , Krste Asanović, Fine-grain CAM-tag cache resizing using miss tags, Proceedings of the 2002 international symposium on Low power electronics and design, August 12-14, 2002, Monterey, California, USA

Parthasarathy Ranganathan , Sarita Adve , Norman P. Jouppi, Reconfigurable caches and their application to media processing, ACM SIGARCH Computer Architecture News, v.28 n.2, p.214-224, May 2000

Christopher J. Hughes , Jayanth Srinivasan , Sarita V. Adve, Saving energy with architectural and frequency adaptations for multimedia applications, Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture, December 01-05, 2001, Austin, Texas

Jun Yang , Youtao Zhang , Rajiv Gupta, Frequent value compression in data caches, Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture, p.258-265, December 2000, Monterey, California, United States

Christopher J. Hughes , Praful Kaul , Sarita V. Adve , Rohit Jain , Chanik Park , Jayanth Srinivasan, Variability in the execution of multimedia applications and implications for architecture, ACM SIGARCH Computer Architecture News, v.29 n.2, p.254-265, May 2001

J. Adam Butts , Gurindar S. Sohi, A static power model for architects, Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture, p.191-201, December 2000, Monterey, California, United States

Michael D. Powell , Amit Agarwal , T. N. Vijaykumar , Babak Falsafi , Kaushik Roy, Reducing set-associative cache energy via way-prediction and selective direct-mapping, Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture, December 01-05, 2001, Austin, Texas

R. Iris Bahar , Srilatha Manne, Power and energy reduction via pipeline balancing, ACM SIGARCH Computer Architecture News, v.29 n.2, p.218-229, May 2001

Dmitry Ponomarev , Gurhan Kucuk , Kanad Ghose, Reducing power requirements of instruction scheduling through dynamic allocation of multiple datapath resources, Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture, December 01-05, 2001, Austin, Texas

I. Kadayif , M. Kandemir , N. Vijaykrishnan , M. J. Irwin , J. Ramanujam, Morphable Cache Architectures: Potential Benefits, ACM SIGPLAN Notices, v.36 n.8, p.128-137, Aug. 2001

Changkyu Kim , Doug Burger , Stephen W. Keckler, An adaptive, non-uniform cache structure for wire-delay dominated on-chip caches, ACM SIGPLAN Notices, v.37 n.10, October 2002

Jinson Koppanalil , Prakash Ramrakhiani , Sameer Desai , Anu Vaidyanathan , Eric Rotenberg, A case for dynamic pipeline scaling, Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems, October 08-11, 2002, Greenoble, France

Gokhan Memik , William H. Mangione-Smith, Increasing power efficiency of multi-core network processors through data filtering, Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems, October 08-11, 2002, Greenoble, France

Ruchira Sasanka , Christopher J. Hughes , Sarita V. Adve, Joint local and global hardware adaptations for energy, ACM SIGOPS Operating Systems Review, v.36 n.5, December 2002

Rajeev Balasubramonian , David Albonesi , Alper Buyuktosunoglu , Sandhya Dwarkadas, Memory hierarchy reconfiguration for energy and performance in general-purpose processor architectures, Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture, p.245-257, December 2000, Monterey, California, United States

Ashutosh S. Dhopakar , James E. Smith, Managing multi-configuration hardware via dynamic working set analysis, ACM SIGARCH Computer Architecture News, v.30 n.2, May 2002

Peter Petrov , Alex Orailoğlu, Towards effective embedded processors in codesigns: customizable partitioned caches, Proceedings of the ninth international symposium on Hardware/software codesign, p.79-84, April 2001, Copenhagen, Denmark

Diana Marculescu , Anoop Iyer, Application-driven processor design exploration for power-performance trade-off analysis, Proceedings of the 2001 IEEE/ACM international conference on Computer-aided design, November 04-08, 2001, San Jose, California

↑ INDEX TERMS

Primary Classification:

B. Hardware

↳ **B.3 MEMORY STRUCTURES**

Additional Classification:

B. Hardware

↳ **B.3 MEMORY STRUCTURES**

General Terms:

Measurement, Performance

↑ Peer to Peer - Readers of this Article have also read:

- Constructing reality

Proceedings of the 11th annual international conference on Systems documentation
Douglas A. Powell , Norman R. Ball , Mansel W. Griffiths

- Data structures for quadtree approximation and compression

Communications of the ACM 28, 9
Hanan Samet

- Toward a real-time Ada design methodology

Proceedings of the conference on TRI-ADA '90
Norman R. Howes

- Fashioning conceptual constructs in Ada

Proceedings of the conference on TRI-ADA '90
Robert C. Shock

- Reuse: the two concurrent life cycles paradigm

Proceedings of the conference on TRI-ADA '90
Richard Drake , William Ett

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



Analytical energy dissipation models for low-power caches

Full text [Pdf \(839 KB\)](#)

Source [International Symposium on Low Power Electronics and Design archive](#)
[Proceedings of the 1997 international symposium on Low power electronics and design table of contents](#)

Monterey, California, United States
 Pages: Pages: 143 - 148
 Year of Publication: 1997
 ISBN:0-89791-903-3

Authors [Milind B. Kamble](#)
[Kanad Ghose](#)

Sponsors [IEEE-CAS : Circuits & Systems](#)
[SIGDA: ACM Special Interest Group on Design Automation](#)

Publisher ACM Press New York, NY, USA

Additional Information: [references](#) [citations](#)

Tools and Actions: [Discussions](#) [Find similar Articles](#) [Review this Article](#)
[Save this Article to a Binder](#) [Display in BibTeX Format](#)

DOI Bookmark: Use this link to bookmark this Article: <http://doi.acm.org/10.1145/263272.263310>
[What is a DOI?](#)

↑ REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

ERB+ 95 [John H. Edmondson](#) , [Paul I. Rubinfeld](#) , [Peter J. Bannon](#) , [Bradley J. Benschneider](#) , [Debra Bernstein](#) , [Ruben W. Castelino](#) , [Elizabeth M. Cooper](#) , [Daniel E. Dever](#) , [Dale R. Donchin](#) , [Timothy C. Fischer](#) , [Anil K. Jain](#) , [Shekhar Mehta](#) , [Jeanne E. Meyer](#) , [Ronald P. Preston](#) , [Vidya Rajagopalan](#) , [Chandrasekhara Somanathan](#) , [Scott A. Taylor](#) , [Gilbert M. Wolrich](#), [Internal organization of the Alpha 21164, a 300-MHz 64-bit quad-issue CMOS RISC microprocessor](#), [Digital Technical Journal](#), v.7 n.1, p.119-135, Jan. 1995

Intel 96 Intel Corporation, Pentium-Pro data book, 1996.

JBD+ 93 3"ouppi, N., et al., "A 300-MHz 115-W 32-b Bipolar ECL Microprocessor", in [IEEE Journal of Solid-State Circuits](#), Nov. 1993, pp. 1152-1165.

KaGh 97 Kamble, M. B. and Ghose, tC, "Energy-Efficiency of VLSI Caches: A Comparative Study", in [Prec. IEEE 10-th, Int'l, Conf. on VLSI Design](#), Jan. 1997, pp. 261-267,

Larus 96 Larus, J., "SHM: A MIPS 2000 Simulator", available form Univ. W'm., CS ftp si~.

Mon 96 Montanaro, J. et al., "A 160 MHz, 32b 0.5 W CMOS RISC Microprocessor", in [IEEE ISSCC 1996 Digest of Papers](#), 1996.

Ro 96 Rogers, A., "CL-SPIM: A Cycle Level Simulator for the MIPS 2000", available from Univ. Wis., CS ftp site,

Smith 82 Alan Jay Smith, Cache Memories, ACM Computing Surveys (CSUR), v.14 n.3, p.473-530, Sept. 1982

StBu 95 Mircea R. Stan, Wayne P. Burleson, Bus-invert coding for low-power I/O, IEEE Transactions on Very Large Scale Integration (VLSI) Systems, v.3 n.1, p.49-58, March 1995

SuDe 95 Ching-Long Su, Alvin M. Despain, Cache design trade-offs for power and performance optimization: a case study, Proceedings 1995 international symposium on Low power design, p.63-68, April 23-26, 1995, Dana Point, California, United States

WeEs 93 Neil H. E. Weste, Kamran Eshraghian, Principles of CMOS VLSI design: a systems perspective, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, 1985

Wijo 94 Wilton, S. E., and 3"ouppi, N., "An Enhanced Access and Cycle Time Model for On-Chip Caches", DEC WRL Research Report 93/5, July 1994

↑ CITINGS 47

Peter Petrov, Alex Orailoglu, Data cache energy minimizations through programmable tag size matching to the applications, Proceedings of the international symposium on Systems synthesis, September 30-October 03, 2001, Montréal, P.Q., Canada

Wen-Tsong Shiue, Chaitali Chakrabarti, Memory exploration for low power, embedded systems, Proceedings of the 36th ACM/IEEE conference on Design automation conference, p.140-145, June 21-25, 1999, New Orleans, Louisiana, United States

Luca Benini, Alberto Macii, Enrico Macii, Massimo Poncino, Synthesis of application-specific memories for power optimization in embedded systems, Proceedings of the 37th conference on Design automation, p.300-303, June 05-09, 2000, Los Angeles, California, United States

G. Esakkimuthu, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, Memory system energy (poster session): influence of hardware-software optimizations, Proceedings of the 2000 international symposium on Low power electronics and design, p.244-246, July 25-27, 2000, Rapallo, Italy

Peter Petrov, Alex Orailoglu, Low-power data memory communication for application-specific embedded processors, Proceedings of the 15th international symposium on System Synthesis, October 02-04, 2002, Kyoto, Japan

Luca Benini, Giovanni De Micheli, System-level power optimization: techniques and tools, Proceedings 1999 international symposium on Low power electronics and design, p.288-293, August 16-17, 1999, San Diego, California, United States

W. Ye, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, The design and use of simplepower: a cycle-accurate energy estimation tool, Proceedings of the 37th conference on Design automation, p.340-345, June 05-09, 2000, Los Angeles, California, United States

Chunho Lee, Johnson Kin, Miodrag Potkonjak, William H. Mangione-Smith, Designing power efficient hypermedia processors, Proceedings 1999 international symposium on Low power electronics and design, p.276-278, August 16-17, 1999, San Diego, California, United States

R. Iris Bahar, Gianluca Albera, Srilatha Manne, Power and performance tradeoffs using various caching strategies, Proceedings 1998 international symposium on Low power electronics and design, p.64-69, August 10-12, 1998, Monterey, California, United States

Michael Powell , Se-Hyun Yang , Babak Falsafi , Kaushik Roy , T. N. Vijaykumar, Gated-Vdd: a circuit technique to reduce leakage in deep-submicron cache memories, Proceedings of the 2000 international symposium on Low power electronics and design, p.90-95, July 25-27, 2000, Rapallo, Italy

Johnson Kin , Munish Gupta , William H. Mangione-Smith, The filter cache: an energy efficient memory structure, Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture, p.184-193, December 01-03, 1997, Research Triangle Park, North Carolina, United States

Koji Inoue , V. G. Moshnyaga , K. Murakami, A history-based I-cache for low-energy multimedia applications, Proceedings of the 2002 international symposium on Low power electronics and design, August 12-14, 2002, Monterey, California, USA

Sari L. Coumeri , Donald E. Thomas, Memory modeling for system synthesis, Proceedings 1998 international symposium on Low power electronics and design, p.179-184, August 10-12, 1998, Monterey, California, United States

Yanbing Li , Jörg Henkel, A framework for estimation and minimizing energy dissipation of embedded HW/SW systems, Proceedings of the 35th annual conference on Design automation conference, p.188-193, June 15-19, 1998, San Francisco, California, United States

Yanbing Li , Jörg Henkel, A framework for estimating and minimizing energy dissipation of embedded HW/SW systems, Readings in hardware/software co-design, Kluwer Academic Publishers, Norwell, MA, 2001

Jouml;rg Henkel , Yanbing Li, Energy-conscious HW/SW-partitioning of embedded systems: a case study on an MPEG-2 encoder, Proceedings of the sixth international workshop on Hardware/software codesign, p.23-27, March 1998, Seattle, Washington, United States

Data memory design and exploration for low-power embedded systems, ACM Transactions on Design Automation of Electronic Systems (TODAES), v.6 n.4, p.553-568, October 2001

I. Kadayif , M. Kandemir , U. Sezer, An integer linear programming based approach for parallelizing applications in On-chip multiprocessors, Proceedings of the 39th conference on Design automation, June 10-14, 2002, New Orleans, Louisiana, USA

S. Kim , N. Vijaykrishnan , M. Kandemir , A. Sivasubramaniam , M. J. Irwin , E. Geethanjali, Power-aware partitioned cache architectures, Proceedings of the 2001 international symposium on Low power electronics and design, p.64-67, August 2001, Huntington Beach, California, United States

William Fornaciari , Donatella Sciuto , Cristina Silvano , Vittorio Zaccaria, A design framework to efficiently explore energy-delay tradeoffs, Proceedings of the ninth international symposium on Hardware/software codesign, p.260-265, April 2001, Copenhagen, Denmark

S. Kim , N. Vijaykrishnan , M. Kandemir , M. J. Irwin, Energy-efficient instruction cache using page-based placement, Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems, November 16-17, 2001, Atlanta, Georgia, USA

Krishna V. Palem , Rodric M. Rabbah , Vincent J. Mooney, III , Pinar Korkmaz , Kiran Puttaswamy, Design space optimization of embedded memory systems via data remapping, ACM SIGPLAN Notices, v.37 n.7, July 2002

Jung-Hoon Lee , Shin-Dug Kim , Charles Weems, Application-adaptive intelligent cache memory system, ACM Transactions on Embedded Computing Systems (TECS), v.1 n.1, p.56-78, November 2002

N. Vijaykrishnan , M. Kandemir , M. J. Irwin , H. S. Kim , W. Ye, Energy-driven integrated hardware-

software optimizations using SimplePower, ACM SIGARCH Computer Architecture News, v.28 n.2, p.95-106, May 2000

H. S. Kim , M. Kandemir , N. Vijaykrishnan , M. J. Irwin, Characterization of memory energy behavior, Workload characterization of emerging computer applications, Kluwer Academic Publishers, Norwell, MA, 2001

Mariagiovanna Sami , Donatella Sciuto , Cristina Silvano , Vittorio Zaccaria, Power exploration for embedded VLIW architectures, Proceedings of the 2000 IEEE/ACM international conference on Computer-aided design, November 05-09, 2000, San Jose, California

Rajeev Balasubramonian , David Albonesi , Alper Buyuktosunoglu , Sandhya Dwarkadas, Memory hierarchy reconfiguration for energy and performance in general-purpose processor architectures, Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture, p.245-257, December 2000, Monterey, California, United States

Bryan Ackland , Chris Nicol, High performance DSPs - what's hot and what's not?, Proceedings 1998 international symposium on Low power electronics and design, p.1-6, August 10-12, 1998, Monterey, California, United States

David H. Albonesi, Selective cache ways: on-demand cache resource allocation, Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture, p.248-259, November 1999, Haifa, Israel

Stefanos Kaxiras , Zhigang Hu , Margaret Martonosi, Cache decay: exploiting generational behavior to reduce cache leakage power, ACM SIGARCH Computer Architecture News, v.29 n.2, p.240-251, May 2001

Zhigang Hu , Stefanos Kaxiras , Margaret Martonosi, Let caches decay: reducing leakage energy via exploitation of cache generational behavior, ACM Transactions on Computer Systems (TOCS), v.20 n.2, p.161-190, May 2002

Luca Benini , Luca Macchiarulo , Alberto Macii , Massimo Poncino, Layout-driven memory synthesis for embedded systems-on-chip, IEEE Transactions on Very Large Scale Integration (VLSI) Systems, v.10 n.2, p.96-105, April 2002

Jayaprakash Pisharath , Alok Choudhary, An integrated approach to reducing power dissipation in memory hierarchies, Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems, October 08-11, 2002, Greenoble, France

David Brooks , Vivek Tiwari , Margaret Martonosi, Wattch: a framework for architectural-level power analysis and optimizations, ACM SIGARCH Computer Architecture News, v.28 n.2, p.83-94, May 2000

Michael Huang , Jose Renau , Seung-Moon Yoo , Josep Torrellas, A framework for dynamic energy efficiency and temperature management, Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture, p.202-213, December 2000, Monterey, California, United States

G. Chen , R. Shetty , M. Kandemir , N. Vijaykrishnan , M. J. Irwin , M. Wolczko, Tuning garbage collection for reducing memory system energy in an embedded java environment, ACM Transactions on Embedded Computing Systems (TECS), v.1 n.1, p.27-55, November 2002

Johnson Kin , Chunho Lee , William H. Mangione-Smith , Miodrag Potkonjak, Power efficient mediaprocessors: design space exploration, Proceedings of the 36th ACM/IEEE conference on Design automation conference, p.321-326, June 21-25, 1999, New Orleans, Louisiana, United States

Daniele Folegnani , Antonio González, Energy-effective issue logic, ACM SIGARCH Computer Architecture News, v.29 n.2, p.230-239, May 2001

Keith I. Farkas , Jason Flinn , Godmar Back , Dirk Grunwald , Jennifer M. Anderson, Quantifying the energy consumption of a pocket computer and a Java virtual machine, ACM SIGMETRICS Performance Evaluation Review, v.28 n.1, p.252-263, June 2000

Timothy Sherwood , Brad Calder , Joel Emer, Reducing cache misses using hardware and software page placement, Proceedings of the 13th international conference on Supercomputing, p.155-164, June 20-25, 1999, Rhodes, Greece

Hojun Shim , Yongsoo Joo , Yongseok Choi , Hyung Gyu Lee , Naehyuck Chang, Low-energy off-chip SDRAM memory systems for embedded applications, ACM Transactions on Embedded Computing Systems (TECS), v.2 n.1, p.98-130, February 2003

Victor Delaluz , Mahmut Kandemir , N. Vijaykrishnan , Anand Sivasubramaniam , Mary Jane Irwin, Hardware and Software Techniques for Controlling DRAM Power Modes, IEEE Transactions on Computers, v.50 n.11, p.1154-1173, November 2001

Luca Benini , Alberto Macii , Massimo Poncino, Energy-aware design of embedded memories: A survey of technologies, architectures, and optimization techniques, ACM Transactions on Embedded Computing Systems (TECS), v.2 n.1, p.5-32, February 2003

Luca Benini , Giovanni de Micheli, System-level power optimization: techniques and tools, ACM Transactions on Design Automation of Electronic Systems (TODAES), v.5 n.2, p.115-192, April 2000

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

Terms used **microinstruction** and **power control** and **cache** and **disable** or **disabled** or **disabling**

Found 4,365 of 113,497

Sort results
by

relevance 

 [Save results to a Binder](#)

[Try an Advanced Search](#)

Display
results

expanded form 

 [Search Tips](#)
 [Open results in a new window](#)

[Try this search in The ACM Guide](#)

Results 61 - 80 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 

61 Specialization tools and techniques for systematic optimization of system software

Dylan McNamee, Jonathan Walpole, Calton Pu, Crispin Cowan, Charles Krasic, Ashvin Goel, Perry Wagle, Charles Consel, Gilles Muller, Renaud Marlet

May 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 2

Full text available:  [pdf\(178.52 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Specialization has been recognized as a powerful technique for optimizing operating systems. However, specialization has not been broadly applied beyond the research community because current techniques based on manual specialization, are time-consuming and error-prone. The goal of the work described in this paper is to help operating system tuners perform specialization more easily. We have built a specialization toolkit that assists the major tasks of specializing operating systems. We de ...

Keywords: operating system specialization, optimization, software architecture

62 Integrating security in a large distributed system

M. Satyanarayanan

August 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 3

Full text available:  [pdf\(2.90 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Andrew is a distributed computing environment that is a synthesis of the personal computing and timesharing paradigms. When mature, it is expected to encompass over 5,000 workstations spanning the Carnegie Mellon University campus. This paper examines the security issues that arise in such an environment and describes the mechanisms that have been developed to address them. These mechanisms include the logical and physical separation of servers and clients, support for secure communication ...

63 Spritely NFS: experiments with cache-consistency protocols

V. Srinivasan, J. Mogul

November 1989 **ACM SIGOPS Operating Systems Review, Proceedings of the twelfth ACM symposium on Operating systems principles**, Volume 23 Issue 5

Full text available:  [pdf\(1.50 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

File caching is essential to good performance in a distributed system, especially as processor speeds and memory sizes continue to improve rapidly while disk latencies do not. Stateless-server systems, such as NFS, cannot properly manage client file caches. Stateful systems, such as Sprite, can use explicit cache consistency protocols to improve both cache

consistency and overall performance. By modifying NFS to use the Sprite cache consistency protocols, we isolate the effects o ...

64 Partitioned instruction cache architecture for energy efficiency

Soontae Kim, N. Vijaykrishnan, Mahmut Kandemir, Anand Sivasubramaniam, Mary Jane Irwin
May 2003 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 2 Issue 2

Full text available: [pdf\(817.81 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The demand for high-performance architectures and powerful battery-operated mobile devices has accentuated the need for low-power systems. In many media and embedded applications, the memory system can consume more than 50% of the overall system energy, making it a ripe candidate for optimization. To address this increasingly important problem, this article studies energy-efficient cache architectures in the memory hierarchy that can have a significant impact on the overall system energy ...

Keywords: Caches, energy, memory system

65 Cache performance for multimedia applications

Nathan T. Slingerland, Alan Jay Smith
June 2001 **Proceedings of the 15th international conference on Supercomputing**

Full text available: [pdf\(642.63 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The caching behavior of multimedia applications has been described as having high instruction reference locality within small loops, very large working sets, and poor data cache performance due to non-locality of data references. Despite this, there is no published research deriving or measuring these qualities. Utilizing the previously developed Berkeley Multimedia Workload, we present the results of execution driven cache simulations with the goal of aiding future media processing architect ...

Keywords: CPU caches, cache, multimedia, simulation, trace driven simulation

66 The shared regions approach to software cache coherence on multiprocessors

Harjinder S. Sandhu, Benjamin Gamsa, Songnian Zhou
July 1993 **ACM SIGPLAN Notices , Proceedings of the Fourth ACM SIGPLAN symposium on Principles & practice of parallel programming**, Volume 28 Issue 7

Full text available: [pdf\(1.12 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The effective management of caches is critical to the performance of applications on shared-memory multiprocessors. In this paper, we discuss a technique for software cache coherence that is based upon the integration of a program-level abstraction for shared data with software cache management. The program-level abstraction, called Shared Regions, explicitly relates synchronization objects with the data they protect. Cache coherence algorithms are presented which use the i ...

67 The Rio file cache: surviving operating system crashes

Peter M. Chen, Wee Teck Ng, Subhachandra Chandra, Christopher Aycock, Gurushankar Rajamani, David Lowell
September 1996 **Proceedings of the seventh international conference on Architectural support for programming languages and operating systems**, Volume 31, 30 Issue 9 , 5

Full text available: [pdf\(1.12 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

One of the fundamental limits to high-performance, high-reliability file systems is memory's vulnerability to system crashes. Because memory is viewed as unsafe, systems periodically write data back to disk. The extra disk traffic lowers performance, and the delay period

before data is safe lowers reliability. The goal of the Rio (RAM I/O) file cache is to make ordinary main memory safe for persistent storage by enabling memory to survive operating system crashes. Reliable memory enables a syste ...

68 Memory hierarchy reconfiguration for energy and performance in general-purpose processor architectures

Rajeev Balasubramonian, David Albonesi, Alper Buyuktosunoglu, Sandhya Dwarkadas
December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture**

Full text available: [!\[\]\(b24c916560b9f61306da39645d78552c_img.jpg\) pdf\(155.56 KB\)](#)
[!\[\]\(9db22c7149b8e973e9e953e59f6c82ab_img.jpg\) ps\(663.39 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



69 Code optimization - I: Compiler optimization-space exploration

Spyridon Triantafyllis, Manish Vachharajani, Neil Vachharajani, David I. August

Full text available:

[!\[\]\(160d9e0b19a821ba9b0563cb27f6b51c_img.jpg\) pdf\(1.19 MB\)](#) [!\[\]\(83ccf0893cfd57bd50e115afafd5c1f7_img.jpg\) ps](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

[Publisher Site](#)

To meet the demands of modern architectures, optimizing compilers must incorporate an ever larger number of increasingly complex transformation algorithms. Since code transformations may often degrade performance or interfere with subsequent transformations, compilers employ predictive heuristics to guide optimizations by predicting their effects a priori. Unfortunately, the unpredictability of optimization interaction and the irregularity of today's wide-issue machines severely limit the accura ...

70 CRL: high-performance all-software distributed shared memory

K. L. Johnson, M. F. Kaashoek, D. A. Wallach
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available: [!\[\]\(f08bfb6107b8810154b81e17d5a88ffd_img.jpg\) pdf\(2.02 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

71 SAFKASI: a security mechanism for language-based systems

Dan S. Wallach, Andrew W. Appel, Edward W. Felten
October 2000 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 9 Issue 4

Full text available: [!\[\]\(796cd277752e45a0cd4157a98fd1f6eb_img.jpg\) pdf\(234.89 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In order to run untrusted code in the same process as trusted code, there must be a mechanism to allow dangerous calls to determine if their caller is authorized to exercise the privilege of using the dangerous routine. Java systems have adopted a technique called stack inspection to address this concern. But its original definition, in terms of searching stack frames, had an unclear relationship to the actual achievement of security, overconstrained the implementation of a Java system, lim ...

Keywords: Internet, Java, WWW, access control, applets, security-passing style, stack inspection

72 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz
December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4

Full text available: [!\[\]\(171b514ff0e24b0f3fcd5e0f946e3e80_img.jpg\) pdf\(5.33 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

73 Fine-grain access control for distributed shared memory

Ioannis Schoinas, Babak Falsafi, Alvin R. Lebeck, Steven K. Reinhardt, James R. Larus, David A. Wood

November 1994 **Proceedings of the sixth international conference on Architectural support for programming languages and operating systems**, Volume 29 , 28 Issue 11 , 5

Full text available: [!\[\]\(a4394173c404dd23dd76a3f299550c41_img.jpg\) pdf\(1.20 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper discusses implementations of fine-grain memory access control, which selectively restricts reads and writes to cache-block-sized memory regions. Fine-grain access control forms the basis of efficient cache-coherent shared memory. This paper focuses on low-cost implementations that require little or no additional hardware. These techniques permit efficient implementation of shared memory on a wide range of parallel systems, thereby providing shared-memory codes with a portability ...

74 Constructing instruction traces from cache-filtered address traces (CITCAT)

Charlton D. Rose, J. Kelly Flanagan

December 1996 **ACM SIGARCH Computer Architecture News**, Volume 24 Issue 5

Full text available: [!\[\]\(597e08ee21112f88d4c9961fafa9d200_img.jpg\) pdf\(595.54 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Instruction traces are useful tools for studying many aspects of computer systems, but they are difficult to gather without perturbing the systems being traced. In the past, researchers have collected instruction traces through various techniques, including single-stepping, instruction inlining, hardware monitoring, and processor simulation. These approaches, however, fail to produce accurate traces because they interfere with the processor's normal execution. Because processors are deterministic ...

75 CHOSarc: kernel support for multiweight objects, invocations, and atomicity in real-time multiprocessor applications

Ahmed Gheith, Karsten Schwan

February 1993 **ACM Transactions on Computer Systems (TOCS)**, Volume 11 Issue 1

Full text available: [!\[\]\(c824d564d2d5d2752eef72d91a8e177a_img.jpg\) pdf\(2.81 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

CHOSarc is an object-based multiprocessor operating system kernel that provides primitives with which programmers may easily construct objects of differing types and object invocations of differing semantics, targeting multiprocessor systems, and real-time applications. The CHOSarc can guarantee desired performance and functionality levels of selected computations in real-time applications. Such guarantees can be made despite poss ...

76 File system usage in Windows NT 4.0

Werner Vogels

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles**, Volume 33 Issue 5

Full text available: [!\[\]\(5d011804813b87f95ca7188a9f39b38c_img.jpg\) pdf\(1.75 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We have performed a study of the usage of the Windows NT File System through long-term kernel tracing. Our goal was to provide a new data point with respect to the 1985 and 1991 trace-based File System studies, to investigate the usage details of the Windows NT file

system architecture, and to study the overall statistical behavior of the usage data. In this paper we report on these issues through a detailed comparison with the older traces, through details on the operational characteristics and ...

77 Improving end-to-end performance of the Web using server volumes and proxy filters

Edith Cohen, Balachander Krishnamurthy, Jennifer Rexford

October 1998 **ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication**, Volume 28 Issue 4

Full text available: [pdf\(1.79 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The rapid growth of the World Wide Web has caused serious performance degradation on the Internet. This paper offers an end-to-end approach to improving Web performance by collectively examining the Web components --- clients, proxies, servers, and the network. Our goal is to reduce user-perceived latency and the number of TCP connections, improve cache coherency and cache replacement, and enable prefetching of resources that are likely to be accessed in the near future. In our scheme, server re ...

Keywords: Web, caching, coherency, filters, piggybacking, prefetching, volumes

78 Removing architectural bottlenecks to the scalability of speculative parallelization

Milos Prvulovic, María Jesús Garzarán, Lawrence Rauchwerger, Josep Torrellas

May 2001 **ACM SIGARCH Computer Architecture News , Proceedings of the 28th annual international symposium on Computer architecture**, Volume 29 Issue 2

Full text available: [pdf\(1.13 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Speculative thread-level parallelization is a promising way to speed up codes that compilers fail to parallelize. While several speculative parallelization schemes have been proposed for different machine sizes and types of codes, the results so far show that it is hard to deliver scalable speedups. Often, the problem is not true dependence violations, but sub-optimal architectural design. Consequently, we attempt to identify and eliminate major architectural bottlenecks that limit the scalability ...

79 Object combining: A new aggressive optimization for object intensive programs

Ronald Veldema, J. H. Cariel, F. H. Rutger, E. Henri

November 2002 **Proceedings of the 2002 joint ACM-ISCOPE conference on Java Grande**

Full text available: [pdf\(99.27 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Object combining tries to put objects together that have roughly the same life times in order to reduce strain on the memory manager and to reduce the number of pointer indirections during a program's execution. Object combining works by appending the fields of one object to another, allowing allocation and freeing of multiple objects with a single heap (de) allocation. Unlike object *inlining*, which will only optimize objects where one has a (unique) pointer to another, our optimization al ...

Keywords: Java, garbage collection, object management

80 Requirements for a layered software architecture supporting cooperative multi-user interaction

Flavio De Paoli, Andrea Sosio

May 1996 **Proceedings of the 18th international conference on Software engineering**

Full text available: [pdf\(1.16 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)